REMARKS ON WARNING SYSTEM SPECIFICATIONS AND STRUCTURES

Robert D. Turner

Institute for Defense Analyses
Arlington, Virginia

August 1969

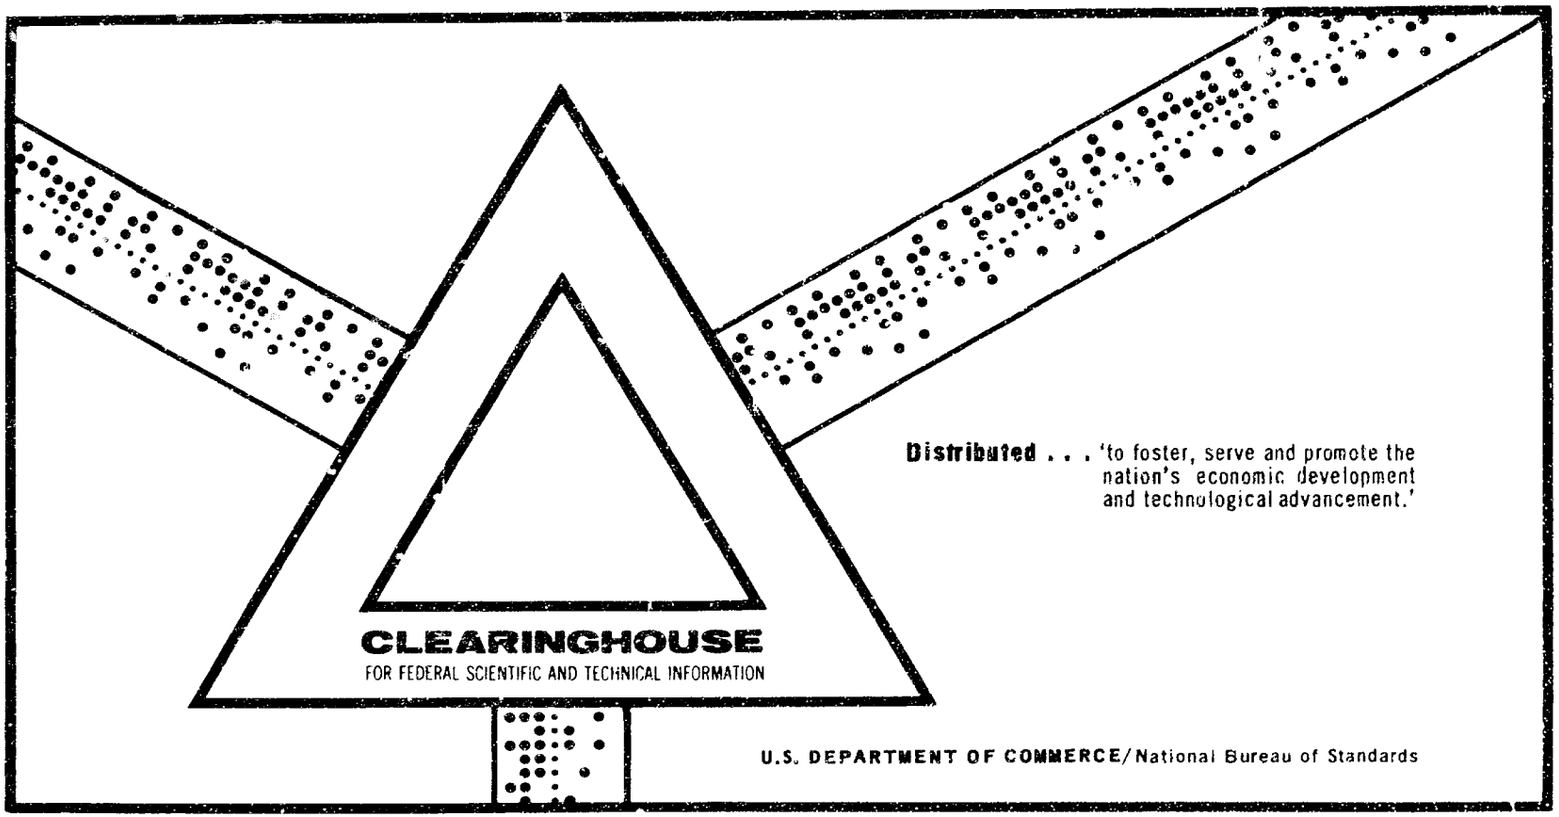

RESEARCH PAPER P-534

REMARKS ON WARNING SYSTEM SPECIFICATIONS AND STRUCTURES

Robert D. Turner

August 1969

INSTITUTE FOR DEFENSE ANALYSES
SCIENCE AND TECHNOLOGY DIVISION

IDA Log No. HQ 69-10632
Copy 79 of 125 copies

58

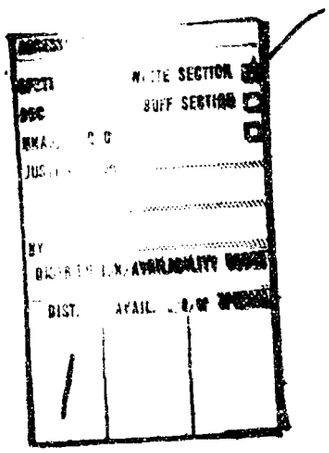

The publication of this IDA Research Paper does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that agency.

This document has been approved for public release and sale; its distribution is unlimited.

RESEARCH PAPER P-534

REMARKS ON WARNING SYSTEM SPECIFICATIONS AND STRUCTURES

Robert D. Turner

August 1969

INSTITUTE FOR DEFENSE ANALYSES
SCIENCE AND TECHNOLOGY DIVISION
400 Army-Navy Drive, Arlington, Virginia 22202

Contract DAHC15 67 C 0011
ARPA Assignment 5

ABSTRACT

Several conceptual questions dealing with warning system structures are discussed, including techniques for correlating the outputs of multiple warning systems that compose a warning complex. The discussions are illustrated by a number of elementary analyses, dealing with warning probabilities, false-warning rates, and warning times. The mathematical calculations pertain to hypothetical automatic event detectors and automatic decision systems employing the outputs of such detectors. It is demonstrated that the use of multiple sensing elements in a warning system implies a need for proper association of redundantly detected signatures of real events and false signatures. The consequences of imperfect association are false-warning rates that can exceed by a considerable factor the rates that would be obtained assuming perfect association and when undesired warnings are generated in response to nonbelligerent activities.

CONTENTS

I. INTRODUCTION

The term "warning system" is used in the present context to refer to the apparatus and operational procedures used to provide the National Command Authority (NCA) with knowledge of an immediately impending attack against the United States and its strategic offensive and defensive forces, e.g., by detection of a large number of Intercontinental Ballistic Missiles (ICBM) launched toward the continental United States. This aspect of the overall strategic forecasting process is usually designated tactical warning, as distinct from strategic warning, which alerts the NCA that such an attack may occur in the near future.

The fundamental differences between tactical and strategic warning are quite important. Strategic warning data are generally prepared for a longer time range and can, within the limits of current policy, only influence the state of preparedness of the nation for general war. More important, strategic warning systems can be exercised in terms of potential conflict at levels lower than general war; the responsiveness and accuracy of the predictions provided by the strategic warning apparatus can thus be evaluated in a variety of situations less critical than general war. Over a period of time, therefore, a good strategic warning capability can, in principle, acquire an image of high credibility in the eyes of the NCA, by providing timely and accurate forecasts of impending crises.

Tactical warning systems cannot be evaluated in the same way; they can only be subjected to limited-scale tests, from which their performance must be inferred, through careful analysis of the test results and thorough understanding of the phenomenology that influences their operation. Such tests are provided by tests of missile forces by other nations, in addition to planned trials and evaluations by the United States.

Once a warning message from a tactical warning system has been received, different responses can be initiated: the commitment of U.S. retaliatory forces, the delegation of authority to use nuclear weapons to defensive forces, and other major military and civilian actions by the nation.

The availability of a credible warning capability could permit substantial savings in the cost of operating the strategic bomber fleet by minimizing the need for airborne alert missions. Other potential benefits have apparently not been evaluated; in general, the value of increased quality of warning data and increased warning time to the NCA is not known. One possibility is the enabling of an option for launch-on-warning (LOW), whereby knowledge of an attack already in being could be used to launch retaliatory weapons. Two kinds of credibility for tactical warning are required, however, for the benefits of LOW to be realized. First, the system or systems composing a tactical warning complex must provide some degree of certainty of detecting the onset of an attack and near certain identification of the attacker; second, the likelihood of declaring that an attack has begun, when such is not the case, must be made vanishingly small. It is well known that these two requirements work against each other.

The tactical warning situation is further complicated by the existence of a variety of strategic threats, other than the ICBM already mentioned. The Submarine-Launched Ballistic Missile (SLBM) threat is one example: the relatively short flight time for such weapons reduces the time available for warning, and the fact that the missile is launched from a mobile platform in the open seas makes more difficult the task of identifying the attacking nation.

The prognosis for realizing a nearly ideal tactical warning capability is not completely bleak, despite the complications and conflicting requirements noted above. Realization is dependent on an understanding of how tactical warning system outputs are to be used; it is the purpose of this paper to indicate some of the statistical and structural concepts that are involved in achieving such understanding. The

substantive text begins with a delineation of design objectives for warning, and then considers the conceptual problems of correlating the outputs of multiple warning systems. These discussions are followed by several rudimentary mathematical analyses pertaining to the performance of (hypothetical) automatic event detectors and decision elements that could constitute a warning capability. The paper is intended to be heuristic; references to specific techniques for acquiring and interpreting tactical warning data have been scrupulously avoided. This approach permits the use of relatively elementary statistical models for the purpose of illustrating concepts that have arisen from studies of real systems. Those participating in the design and evaluation of real systems must, of course, ascertain the real statistical descriptions.

II. THE PROBLEM

A. PRINCIPAL OBJECTIVES

Several of the principal objectives in synthesizing a ta warning capability are discussed in the following paragraphs.

1. Adequate Spatial Coverage means that all points on t of the earth, all regions of the atmosphere, and all regions from which an attack could be launched should be kept under s lance for warning purposes. The qualification "adequate" int with the reaction time objective noted below in point 5. For it may be unnecessary to maintain surveillance of launch poin bomber attacks, because the time available for dealing with su attack (when detected several hundred miles from the United St long compared with the time available for dealing with other l attacks. Thus, for bomber attack warning, the surveillance re could be strongly curtailed, relative to the surveillance req quired for timely detection of other kinds of attacks.

2. Adequate Temporal Coverage simply means that the perf of the total tactical warning capability should be insensitive warning phenomena that may influence the performance of indivi tem elements.

3. High Probability of Attack Detection is influenced by tent to which objectives 1 and 2 are achieved, but it also inv quality of the detection equipment used and the means employed translating detections into warning messages. The qualificati does not mean arbitrarily close to unity, because the warning is but one element in an overall deterrent capability. It is all obvious, for example, that greater deterrence would be ach

increasing the probability of attack detection from, say, 0.95 to 0.99. In either case, the odds against achieving an attack without tactical warning to the United States deterrent complex are formidable. A decision to attack the United States in the face of such odds would most likely still be made even if the attack warning probability were unity.

4. High Confidence in Attacker Designation means that in addition to declaring with high confidence that the United States is under attack, the warning capability must provide the command-and-control structure with the identity (or identities) of the attacking nation (or nations). This is essential if the doctrine of retaliation is to be implemented. The objective also clearly involves objectives 1, 2, and 3. In many instances, designation of the locations of launch points is sufficient for this purpose, and such data would be a natural output of several types of warning systems. It should be noted, however, that it is possible to postulate irrational threats for which no practical solution to the attacker designation problem has been delineated. This point evokes the notion of an interaction between tactical warning and defense; if the nature of an attack is such that a high-confidence designation of the attacking nation is impossible, then retaliation against such a threat may be ruled out, the deterrent value of retaliation is nullified, and defense is the only available option.

5. Short Reaction Time can equivalently be called maximum advance notice. It will be seen that the reaction time associated with a tactical warning system depends on the kinds and quality of data that it provides and on the character of the threat that is being detected; clearly the need for advance notice depends on the use that will be made of the warning data. There is obviously an upper limit to the advance notice available, ranging from several hours for bomber attacks to several minutes for attacks by submarine-launched ballistic missiles (SLBM). For certain future threat types, the potentially available advance notice cannot be predetermined, but it may be greater than the notice available for ICBM attacks or less than that available for SLBM attacks, depending on the capabilities of the warning system and the definition of tactical warning. For these threats, the distinction

between strategic warning and tactical warning becomes vague, which is not to say that achievement of a technical capability for inferring the existence of a threat situation is unimportant.

Reaction time requirements are not easy to specify, since they depend on the extent to which prior alerting by means of strategic warning capabilities can be achieved and the use that will be made of the tactical warning data by the NCA.

6. Low False-Warning Rate simply means that the rate (e.g., number of times per year or decade) at which a system generates warning messages when an attack is not occurring shall be tolerably low, the tolerance level being established by the implications of the response to the message. This suggests that warning messages may be categorized according to the level of the presumed threat; a low-level warning message would generally lead to a low-level response, e.g., alerting of strategic forces. An intermediate-level warning message would presumably lead to such actions as scrambling of strategic bomber forces; only the highest level warning message would trigger a LOW retaliation. Given a spectrum of possible responses to warning messages, there will be a spectrum of tolerable false-warning rates, determined somewhat subjectively by the tolerable cost of false responses.

The false-warning rate specification question is further complicated when deployment of an active defense is undertaken. The availability of defense provides additional options in the response to warning messages and, generally speaking, should permit substantial reductions in higher level false-warning rates, because the thresholds for generating warning messages that lead to higher level responses can be increased, by virtue of the defense option.

B. SECONDARY OBJECTIVES

Realization of the preceding objectives implies the existence of a useful warning capability; the following objectives support the need to sustain that utility.

1. Low Susceptibility to Sensor-Oriented Interference refers to two main occurrences: the credibility of warning messages can be degraded by spoofing (generation of false signatures that are detected by a warning system and interpreted as threatening events), and attack activity signatures can be masked by jamming. Just as a highly credible warning capability is strongly leveraged, so is a capability for degrading its utility, although deliberate attempts at such degradation can in themselves be regarded as strategic warning indicators. A related secondary objective is achievement of insensitivity to other measures intended to degrade warning capabilities (e.g., jamming of communication links or destruction of processing facilities). This is not cited as a primary objective, because it is common to all elements of the command-and-control complex and is dealt with by prudent design practice, including hardening and redundancy.

2. Growth Potential or Residual Capacity for dealing with future threats is important, since the dimensions of possible threats are continually increasing with innovations and advances in offensive weapon technology. Indeed, it can be argued that the development of new enemy offensive weapon systems may be motivated in part by an effort to circumvent the leverage exerted by an effective warning system.

Implicit in all of these objectives is a fundamental requirement that the warning system be capable of determining not only that an attack is occurring, but that the attack is indeed directed toward the United States. Systems that indicate only that an attack is being launched (without specifying the object of the attack) can provide valuable alerting and corroborative functions, but they do not constitute solutions to the warning problem.

C. THE WARNING COMPLEX

Achievement of these objectives seems to imply the development and deployment of several warning systems that are different in kind as well as in coverage; collectively, we refer to these systems as a warning complex. We shall be concerned in what follows with the statistical characterization of individual warning system outputs, the

possible structures of warning complexes, and the statistical characterization of the outputs of a warning complex. The emphasis will be on systems designed to detect ICBM attacks against the United States, centering on false-warning statistics and attack-detection (true warning) probabilities; some consideration will also be given to response-time characteristics of warning complexes. No consideration will be given to the questions of spatial and temporal coverage, although some perhaps unappreciated implications of redundant coverage will be discussed. Finally, we will do no more than allude to the important problem of providing communication links between the sensing elements of the warning systems, the data processing and interpretation elements of the warning complex, and the NCA.

III. WARNING SYSTEMS STRUCTURES AND SPECIFICATIONS

A. THE ELEMENTARY STRUCTURE

Figure 1 depicts an elementary warning system structure. The system is equipped with an array of sensors; each sensor is designed to detect some physical manifestation of an attack or one or more of the elements of an attack. Multiple sensors are generally required, so as to achieve adequate spatial coverage, increased reliability, increased probability of detection, or additional data for subsequent description of the attack.

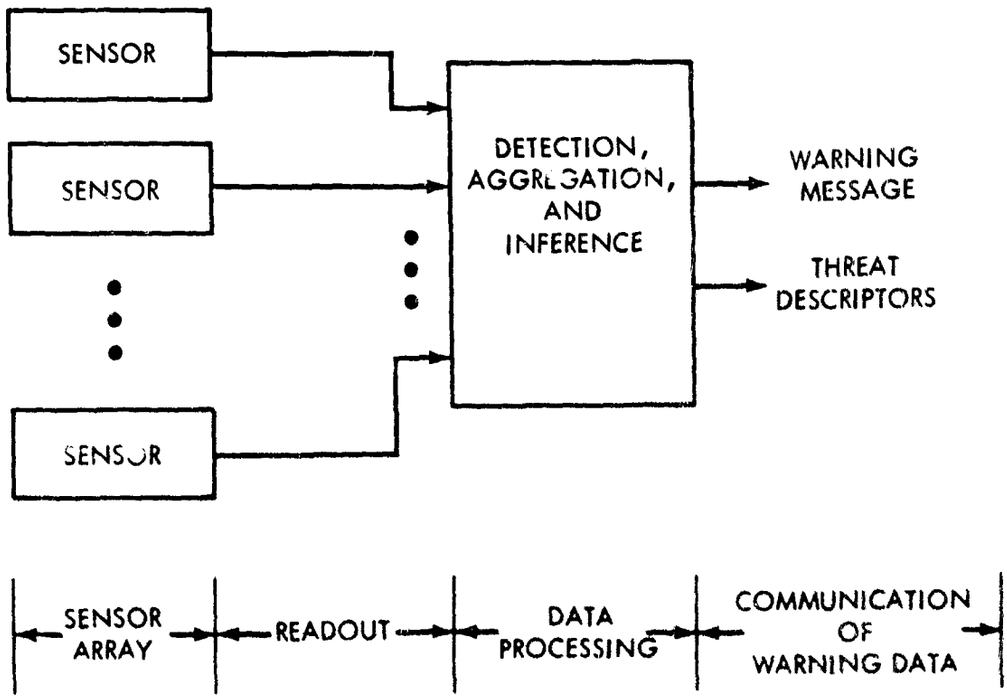

FIGURE 1. Elementary Warning System Structure

The attack or element(s) of an attack that are to be detected will be referred to as an event, and the response of a sensor to an event will be referred to as its signature. In addition to providing a basis for detecting the event, the signature will in general depend on the character of the event (e.g., number of elements in the attack, point of origin). In general, the signature will not explicitly reveal the information that is of interest to the warning system user.

The spatial and temporal coverage provided by the individual sensors may overlap in various ways. Two or more sensors in the array may respond simultaneously to an event, or in some time sequence. The order and time spacing of sequential responses may not be known in advance, in which case the time sequence itself provides information relevant to the event description. The spatial coverage provided by a sensor may be a stochastic process.

It will be assumed that the sensors in the array are functionally identical, in the sense that they are designed to exploit a single common class of physical phenomena, producing real or false signatures under similar conditions.

The outputs of the sensors are delivered (read out) by some means to a data-processing facility, which renders detection decisions, extracts descriptions of detected events, and renders decisions to transmit warning messages. Of particular interest in the descriptions that are inferred are data pertaining to the size of the attack, its point of origin, and parameters for estimating the potential target of the attack. The decision to transmit a warning message is based on the attack-size estimate and an inference from the event-description parameters that the United States is the potential target.

An important function of the data-processing facility is aggregation, which is the process of associating multiple signatures from a single event and compiling a composite multiple-sensor signature (e.g., one including the time-sequence/sensor-number data mentioned previously). The association operation is essential to prevent multiple signature detections (from a single event) from being interpreted as multiple

events. The composite signature may lead to a more precise description of the event and may contain data relevant to the event description that are not available from a single signature.

Finally, the warning message and aggregated event descriptions are transmitted to the user via a communications link. It is assumed that the link is perfect, although certain losses can be accounted for by modifying the parameters characterizing the individual warning systems.

Many ramifications of the foregoing structure can be described; a few will be mentioned here. The system may be equipped with a secondary sensor array, whose function may be to ascertain the operational status of the primary sensors (e.g., coverage) or to provide data for modification of the signature-detection algorithms used in the data-processing facility.

The read-out process may be intermittent (instead of continuous, as was implied above). For example, each sensor may be equipped with a predetection decision element to limit read-out transmissions to those sensed signal sequences that are most likely to contain signatures (autonomous reporting). Alternatively, the sensors may contain storage elements; read-out can then be accomplished on interrogation from the data-processing facility (command reporting). Such interrogations may be transmitted to the sensors cyclically, on the basis of signatures received from other sensors in the array, or on the basis of data from other warning systems.

In addition to interrogation, the data-processing facility may be equipped to command parameter changes in the individual sensors. The read-out process may be accomplished through intermediate collection-and-relay centers, to facilitate reception of data from remote sensors.

These ramifications are generally undertaken to ease the system design (even though they may appear to increase its complexity) or to improve system performance. Very little is known regarding the dynamically optimum control of such a structure.

B. COMPOSITE STRUCTURES

Some rudimentary ideas will be presented here for combining warning systems into a warning complex, in order to establish a conceptual framework for the analytical considerations of Section V. As was noted in Section II, the use of a number of diverse warning systems seems necessary in order to meet the requirements of spatial and temporal coverage, low false-warning rate, high probability of generating a warning message when an attack against the United States is in progress, and so forth. Implicit in these requirements is the ability to monitor activity in all threat classes that the user deems to be significant. The functions of a warning system are to detect the occurrence of events that are characteristic of one or more threat classes and, by analyzing their signatures, to determine whether a threat indeed exists. Because of the diversity of threat classes (ICBM, SLBM, manned bombers, fractional orbital bombardment system (FOBS), and others) and because each threat class presents different kinds of signatures for detection and analysis, the requirement to cover all significant threat classes additionally implies the use of multiple and diverse warning systems.

The general scheme for a warning complex is that each warning system transmits its warning messages and threat descriptions to a central data-processing facility. For at least two reasons, the final responsibility for determining the existence and character of an attack against the United States rests with this facility. First, a not unreasonable attack plan could employ a wide variety of threat classes with a relatively low level of activity in each class. Thus, while a warning system monitoring activity within a single threat class might ascertain the existence of a threat, it would be incapable of assessing the significance of the totality of activity in all threat classes. Second, the need for high credibility in the warning messages that are delivered to the NCA necessitates a capability for verification, at least for the more important threat classes.

The functions of the central data-processing facility are therefore twofold: to integrate messages and descriptions from multiple

warning systems pertaining to each threat class (thereby verifying the existence of threats and obtaining more precise information on them) and to generate for the NCA a composite description of the total attack, including all threat classes that are monitored by the individual systems. The following are some of the complications that must be dealt with in performing these functions:

1. Different warning systems that monitor activity within a single threat class may provide event detections and threat descriptions for that class at grossly different times.
2. The occurrences of false-event detections by different systems may not be statistically independent.
3. Different systems that monitor activity within a single threat class may provide different kinds of descriptions of events that are detected.
4. Discrepancies can exist between the event detections and descriptions provided by different systems that monitor activity within a single threat class.

These problems, and others that will be mentioned, suggest that several different kinds of procedures will be required in the central data-processing facility. Some of these procedures will now be discussed. Before doing this, however, it will be noted that the simple scheme that has been described does not indicate many features that may be required in a real complex, e.g., means for modifying the operation of a component system in response to strategic intelligence or to information received from another system, or means for dynamically allocating communication channel capacity to the several reporting links in response to actions taken against the complex itself.

The first kind of integration that will be considered involves outputs from two similar systems that have a common spatial region under surveillance and that provide nearly coincident event detections and threat descriptions on a common threat class. In processing these reports, the central data-processing facility attempts to match the two sets of descriptions for threats that are detected in the common

surveillance region. This can be done in near-real time, because (by assumption) the detection reports from each system will be nearly coincident. The matching operation is accomplished by comparing the event descriptions* reported by one system with those reported by the other, but only for those events whose descriptions indicate that they have occurred within the common surveillance region for the two systems. The results of this processing can take the following forms:

1. Unambiguously verified events
2. Partially verified events
3. Unresolved (ambiguously verified) events
4. Unverified events
5. Unverifiable events

In the first instance, a single event description from one system matches sufficiently well with a single event description from the other system that the comparison criteria are satisfied; the two sets of data can then be merged into a single report and a single composite event description. In the second instance, the comparison process indicates that the two event descriptions partially agree, but discrepancies disallow high confidence that they can be merged into a single report. Unresolved event verifications result when two or more descriptions from one system agree sufficiently well with one or more descriptions from another system to satisfy the comparison criteria; thus, while the event occurrences are verified, the pairing to form composite descriptions is ambiguous. The fourth case refers to events detected by one system that should have been detected by the other; such cases arise because of failure to detect or because errors in the descriptions cause a failure to match the reports. Unverifiable events are those that are detected outside the common surveillance region and must therefore be treated as reports from an isolated system.

*The most obvious data for purposes of comparison are time of occurrence and kinematic descriptors. Other signature data may be available for comparison purposes, depending on the nature of the sensors. Some data may be useful for verification only through the process of consistency determination, which will be discussed subsequently.

There is a fairly obvious trade-off between the statistics of the output of the matching operation, the stringency of the comparison criteria, and the quality of the event description data. Use of more stringent comparison criteria will lead to fewer false verifications and fewer ambiguous verifications, but to more partially verified and unverified events. The interpretation of the outputs hinges on quantitative knowledge of the joint event-detection and description-error statistics for the two systems.

The next situation to be discussed is that of two similar systems that monitor activity in a common threat class, but provide reports on such activity at different times. In such cases, the descriptions provided by the later reporting system will be time-transformations of the descriptions provided by the earlier reporting system. In this case, the verification process involves predicting, from the earlier descriptions, the quantitative character of the descriptions that will be provided later. An example of this procedure is the use of trajectory data from a boost-phase ICBM detection system to predict the kinematic descriptions that would be observed with a mid-course or early reentry detection system. The output of the integration process has the same structure as was described previously; in general, however, the comparison criteria must be less stringent, to accommodate prediction errors.

The third case to be discussed is the most difficult from the standpoint of data processing. Here, two systems monitor activity in the same threat class, but are so dissimilar that credible verification by direct or predictive description comparison is not possible. Even so, the descriptions provided may have substantially similar content; the question arises as to whether pairings can be made on the basis of a determination of mutual consistency of the data. While a true verification may not be possible, it may be feasible to establish "most likely" associations of the events reported by one system with those reported by another. Consistency determination is accomplished by ascertaining whether there exists a reasonable hypothetical characterization of the event that can lead to a synthetic composite description,

some of whose components agree sufficiently well with the d ption provided by one system, and some which agree sufficiently w with the other description. The problem here is to generate a h hesis set that is sufficiently detailed that it provides the requ s syn-thetic composite description for comparison purposes, but s iciently concise that it allows the central data processor to scan t ntire set (in real time) before rejecting a pairing. To some exte varion techniques for ordering of the search progra can expedite t proce-dure, but the requirements for data processing may still be midable

The final case is one for which integration, as such, i ot pos-sible: when two systems are monitoring activity in disjoint reat classes. Here the data can be aggregated only at the grosse level, and the evaluation problem is one of determining whether the reat manifestations reported have apparent objectives that are cc tent with a reasonable hypothetical attack plan.

These discussions are intended to indicate the complexi f the task of integrating warning data from multiple warning syste The purpose of attempting such a formidable undertaking is of co to provide the NCA with the most accurate and credible descript of an impending attack, despite the diversity of forms that such a tack can assume.

C. SPECIFICATIONS

A subset of possible warning system specifications will dis-cussed here; the example used is that of a warning system de to detect the initiation (launching) of an ICBM attack against t United States. The numbers used to illustrate the specification co s are largely hypothetical, but the concepts are not.

The specification task is complicated by the following f rs:

1. Certain natural phenomena may give rise to signature ected by the system that it interprets as being characteri of missile-launch activity; this contingency will be re to as a natural launch-detection false alarm.*

*Not to be confused with a false warning.

2. The natural false alarms alluded to can be augmented by detections of nonbelligerent launch activity, e.g., launching of nonbelligerent satellite vehicles and tests of missile systems.
3. Under certain conditions, interactions between natural phenomena and nonbelligerent activities can cause the response of a sensor (the signature) to be confused, resulting in misclassification of the event and other errors.

In addition, certain attack-staging procedures may permit an attacker to achieve a sizeable total attack capability without exceeding alerting thresholds that are based on elementary launch-rate (or, more properly, launches per specified unit of time) criteria.

For the moment, an attack will be defined as the launching of N_A or more missiles against United States targets within T_A minutes; it will be required that the system provide at least T_W minutes warning,* with a warning probability not less than P_W. Representative values might be as follows:

$$N_A = 20 \quad (1)$$

$$T_A = 5 \text{ minutes} \quad (2)$$

$$T_W = 20 \text{ minutes} \quad (3)$$

$$P_W = 0.95 \quad (4)$$

The reason for specifying 0.95 for warning probability is beyond the scope of this paper and involves the subjective determination of the contribution made to deterrence by the warning system. The most "deterrent" warning system would provide warning with certainty, but it seems apparent that if the odds against sneaking through the warning system are very high (19-to-1 for the present specification), then the

*As stated here, the specification implies that a warning message must be delivered to the user at least T_W minutes prior to first impact. An alternative requirement is that warning be provided no more than T_W minutes after the first launch. We use the first form because it seems more relevant.

warning system is practically as deterrent as one for which the warning is unity. It remains an open question as to whether a value of 0.8, for example, would be sufficient.

In addition to the warning statistics, the false-warning performance must also be specified. Determination of an acceptable false-warning rate, R_W (the frequency with which the system delivers warning messages to the user when no attack is actually occurring), depends critically on the use that would be made of the warning message. It is to be noted that there would most likely be a spectrum of responses, some of which would be made in response to warning messages generated with different values of N_A, and some of which would be made only in response to verified messages, the verification or corroboration being made by two or more independent warning systems in the complex. For the sake of discussion, it will be assumed that the launch-count criterion specified above results in messages that are used only to alert strategic forces* and are subject to corroboration prior to more substantive commitments. With this restriction, it seems reasonable to assume that

$$R_W = 2 \text{ per year} \tag{5}$$

would be acceptable.

The reaction of the system to the nonbelligerent activity referred to in 2 above must also be specified. We will state this specification somewhat arbitrarily: the probability that the system will deliver a warning message in response to the launching of M or fewer vehicles within an interval of T_A minutes shall not exceed Q_W; for the sake of discussion, the values

$$Q_W = 0.1 \tag{6}$$

$$M = 5 \tag{7}$$

will be adopted.

---

*This need not rule out the possibility of a more stringent criterion, e.g., $N_A' = 100$, for warning messages that would produce a more serious response.

We will not attempt to describe specifications for a warning complex. Some notions as to the form that these might take are implicit in the preceding discussion and in the analyses that follow.

IV. SINGLE-SYSTEM IMPLICATIONS

A. EVENT-DETECTION CAPABILITIES AND WARNING THRESHOLDS

It will be assumed that the decision procedure employed by the warning system in generating warning messages is equivalent to the scheme depicted in Fig. 2. Event detections (e.g., detections of ICBM launches) are delivered to a delay-and-counter combination, which simply counts the number of detections that have occurred in the past T_A minutes. The warning message generator compares the counter reading with a threshold number, N_W, and generates a warning message only when the counter reading goes from N_W-1 to N_W.

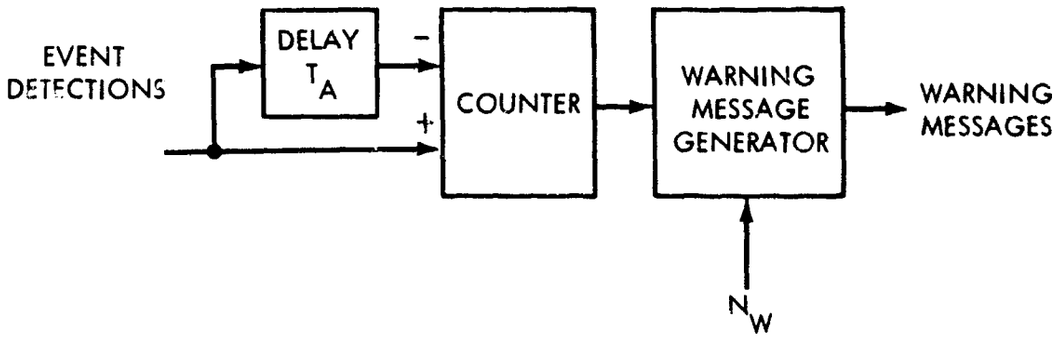

FIGURE 2. Elementary Warning Decision Process

This scheme is applicable regardless of whether the event detections have been previously sifted (on the basis of the corresponding event descriptions) to discriminate against apparent detections whose descriptions do not correspond to known threat classes. It is also entirely conceivable that the warning message generator may employ more than one threshold, generating messages of greater significance when higher counting thresholds are reached.

Now suppose that detections of real events are statistically characterized by the binomial distribution. That is, let P_E denote the probability that a single event is detected, and make two assumptions:

1. A single event leads to a single event detection, at most.
2. Individual event detections are mutually statistically independent.

Then the probability that the launching of N_A missiles within T_A minutes will result in the generation of a warning message is given by

$$P_W = \sum_{n=N_W}^{N_A} \binom{N_A}{n} P_E^n \left(1-P_E\right)^{N_A-n} \qquad (8)$$

Some additional qualifying remarks are in order. First, Eq. 8 does not include the effect of coincidental false event detections that could increase the apparent number of events detected during the onset of the attack; it is therefore conservative in this regard. Second, the occurrence of more than N_A events within T_A minutes will increase the probability of generating a warning message above the value computed using Eq. 8. Third, if the N_A events are spread out over an interval of duration greater than T_A minutes, the probability of generating a warning message will be reduced. In this instance, a (possibly quite crude) lower bound on the probability of generating a warning message is obtained by replacing N_A in Eq. 8 with the maximum number of events occurring within a T_A - minute interval. Finally, the underlying model does not apply to situations for which eventual detection of individual events is practically certain but in which the most significant elements of variability are the times at which such detections take place.

More recondite formulas can be developed that overcome some of these limitations, but Eq. 8 will suffice for the present purpose, which is to delineate a relationship between N_W and P_E in attaining a required value for P_W. Such a relationship is exhibited in Table 1,

which lists the values of P_E required to achieve specified values of P_W for a given threshold N_W, given that $N_A = 20$ events that have occurred within the counting interval T_A.

TABLE 1. EVENT-DETECTION PROBABILITIES (P_E) REQUIRED TO ACHIEVE A SPECIFIED WARNING PROBABILITY OF AN ATTACK BY 20 MISSILES (BINOMIAL MODEL)

Warning Threshold N_W	Required Value of The Probability of Warning P_W			
	0.95	0.90	0.85	0.80
6	0.456	0.415	0.388	0.366
7	0.508	0.467	0.440	0.418
8	0.558	0.519	0.491	0.469
9	0.606	0.567	0.541	0.519
10	0.653	0.615	0.589	0.568
11	0.698	0.662	0.637	0.616
12	0.741	0.707	0.683	0.663
13	0.783	0.751	0.728	0.709
14	0.823	0.793	0.772	0.755
15	0.860	0.834	0.815	0.799
16	0.896	0.873	0.856	0.842
17	0.929	0.910	0.895	0.883
18	0.958	0.944	0.932	0.922
19	0.982	0.973	0.966	0.959
20	0.997	0.995	0.992	0.989

Figure 3 presents some of the data of Table 1 in graphic form and shows that the required values of P_E are relatively insensitive to the specification of P_W, over the range of values being considered. In actual fact, P_E will itself be a variable, depending on the time of occurrence of the event, the location of the event within the coverage domain of the system, and the particular kind of event within the

threat class being monitored. The value of $P_E = 0.70$ will be adopted as representative; Table 1 then indicates that a threshold value $N_W = 11$ can be employed.

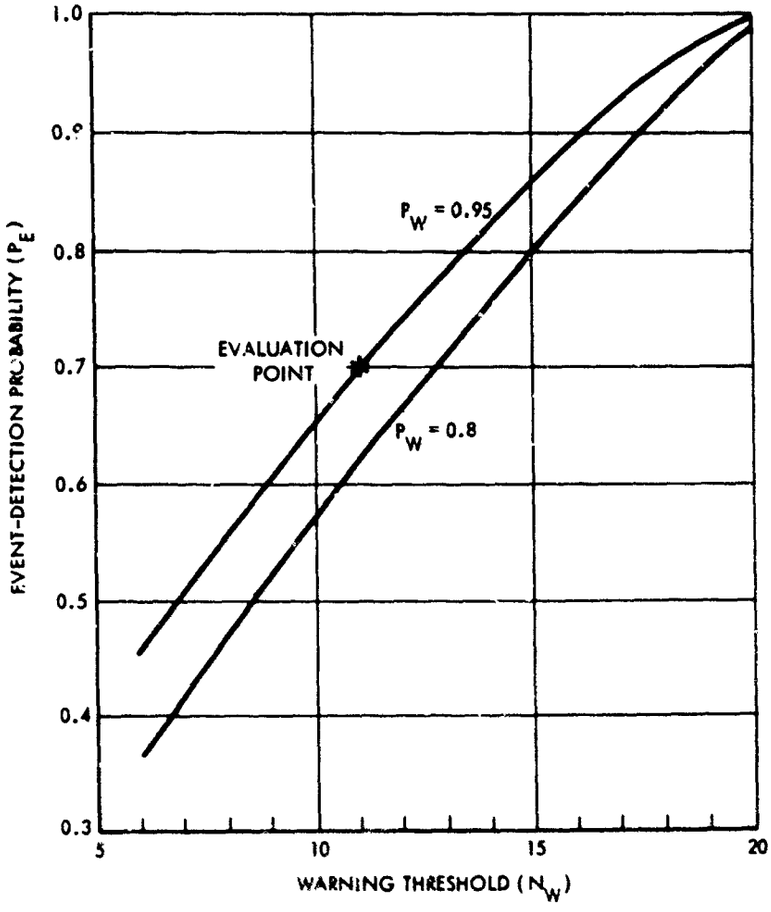

FIGURE 3. Relationship Between Event-Detection Probability, Probability of Warning and Warning Threshold

It is, of course, generally desirable to use as high a threshold value as possible to minimize the false-warning rate, but there are important qualifications to this remark. If the warning system were improved to yield $P_E = 0.99$, for example, an obvious reaction would be to increase the threshold N_W to 19, say. Doing so, however, would increase the possible level of a "sneak-through" attack from 10 launches (in the case $P_W = 0.70$) to 18 launches. By setting the threshold at the lowest possible value, consistent with the tolerable false-warning rate, the range of possible sneak-through attacks is minimized, and the highest possible warning probability for a given attack size is generally obtained.

B. FALSE-DETECTION AND FALSE-WARNING STATISTICS

The average false-warning rate associated with the scheme of Fig. 3 is simply the false-event-detection rate multiplied by the probability that a false-event detection is preceded by exactly $N_W - 1$ event detections in the previous T_A minutes. Assuming that the occurrence of false-event detections is characterized by the Poisson distribution and that the false-event-detection rate is constant, the false-warning rate is given by

$$R_W = R_E \frac{(R_E T_A)^{N_W-1}}{(N_W - 1)!} \exp(-R_E T_A) \qquad (9)$$

where R_E denotes the false-event-detection rate. If R_E is not a constant, but varies in time, then R_W is also dependent on time, and is given by

$$R_W(t) = R_E(t) \frac{[R_E^*(t) T_A]^{N_W-1}}{(N_W - 1)!} \exp[-R_E^*(t) T_A] \qquad (10)$$

where

$$R_E^*(t) = (1/T_A) \int_{t-T_A}^{t} R_E(t') dt' \qquad (11)$$

Eqs. 10 and 11 are presented as a matter of record; in what follows, it will be assumed that R_E is constant.

Table 2 lists some values for R_W for given values of R_E, for the standard case T_A = 5 minutes, N_W = 11. It can be seen that the false-warning rate is quite sensitive to the false-event-detection rate. For a false-warning rate of 2 per year, the tolerable false-event-detection rate is 20.6 per hour, but an increase of R_E to 21.5 per hour will result in R_W being 3 per year.

TABLE 2. FALSE-WARNING RATE AS A FUNCTION OF FALSE-EVENT-DETECTION RATE

False-Event-Detection Rate, R_E (per hour)	False-Warning Rate, R_W (per year)
1	3.6×10^{-14}
2	6.8×10^{-11}
5	1.3×10^{-6}
10	1.7×10^{-3}
20	1.5
50	3.0×10^{3}
100	9.4×10^{4}

Integration Time, T_A, = 5 minutes

Warning Threshold, N_W, = 11

These calculations are valid only if the underlying model for false-event occurrences is satisfactorily approximated by the Poisson distribution. The consequences of one possible departure from simple Poisson statistics will now be discussed; the bulk of the analysis supporting this discussion is given in Appendix A.

It will be assumed that the departure from the simple Poisson model used above can be represented by the superposition of three statistically independent false-event-detection sequences:

1. Singlets: a Poisson-distributed random sequence of false-event detections similar to the sequence implicit in Eq. 9.

2. Doublets: a Poisson-distributed random sequence of pairs of false-event detections, each pair occurring within a time interval that is small compared to the integration time, T_A.

3. Triplets: a Poisson-distributed random sequence of groups of three false-event detections, each group again occurring within a time interval that is small compared to T_A.

Letting $R_E^{(1)}$ denote the average singlet rate, $R_E^{(2)}$ denote the average rate of occurrence of doublets, and $R_E^{(3)}$ denote the average rate of occurrence of triplets, the overall false-event-detection rate is given by

$$R_E = R_E^{(1)} + 2R_E^{(2)} + 3R_E^{(3)} \qquad (12)$$

This model has been assumed because it is more or less analytically tractable; there is a possible mechanism that would support its relevance to real situations.

Suppose that there exist natural phenomena that can give rise to signatures that are interpreted as being real events, and that such signatures are detected in the outputs of more than one sensor in the sensor array of the warning system. If the responses of the sensors to a Poisson-distributed sequence of such natural events are not always perfectly associated, the multiple signatures that are delivered to the system central processor will, on occasion, be interpreted as multiple events. A natural event that produces detectable signatures in two or more sensors in the array can lead to singlets only if the association of the signatures is perfect; it can produce doublets, triplets, or even higher order events if two, three, or more of the signatures are not properly associated.

Figure 4 depicts the results of calculations of R_W, assuming that the overall false-event-detection rate, R_E, is held at 20.6 per hour, which (as was noted previously) yields R_W = 2 per year under the

singlet-only model implicit in Eq. 9. Two cases are shown:

$$R_E^{(3)} = 0.1\, R_E^{(2)}$$

and

$$R_E^{(3)} = 0.01\, R_E^{(2)}$$

As before, T_A = 5 minutes and N_W = 11. It can be seen that the false-warning rate is quite sensitive to the occurrence of doublets and triplets, approximately doubling when $R_E^{(2)}$ = 0.5 per hour; this is roughly equivalent to having a 97.5 percent probability of satisfactory association of two signatures from a single natural event. When $R_E^{(2)}$ = 1 per hour, the singlet rate is about 18.3 per hour; approximately 10 percent of the false-event detections appear as doublets and triplets. The false-warning rate is then 8 per year for $R_E^{(3)}$ = 0.01 per hour and 12 per year for $R_E^{(3)}$ = 0.1 per hour. For an example of the consequences of rather poor association, take $R_E^{(1)}$ = 6.8 per hour, $R_E^{(2)}$ = 6 per hour, and $R_E^{(3)}$ = 0.6 per hour; the average false-warning rate is then approximately 180 per year, or about one every 2 days.

As an extreme example, suppose that all such false-event detections appear as doublets, with an average occurrence rate of 10.3 per hour (or 20.6 false-event detections per hour); then the analysis (Appendix A) shows that the average false-warning rate is 149 per year. Comparing this with the 180-per-year figure just cited shows the extreme sensitivity to triplets and higher order occurrences. If all the 20.6 per-hour false-event detections appear as triplets, the resulting false-warning rate is over 1000 per year (approximately 3 per day).

To summarize, these results indicate that the false-warning performance of a warning system is critically dependent on the temporal statistical structure of the events that can cause false warnings. In particular, it would seem that great care must be exercised to provide very high probabilities of perfect association of multiple responses to natural events.

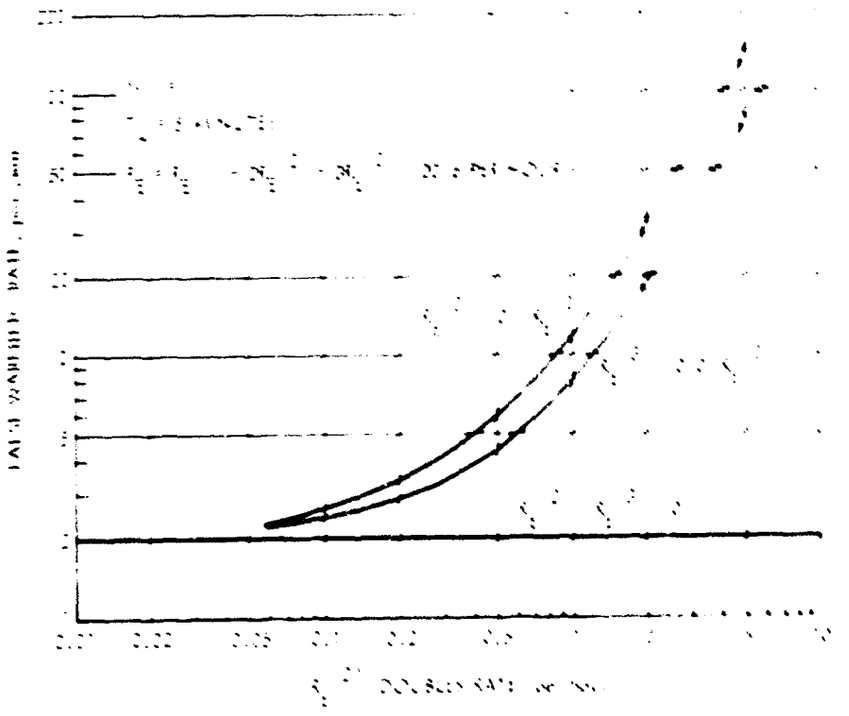

FIGURE 4. False-Warning Rate for Compound False-Event Detections

C. REACTIONS TO NONBELLIGERENT ACTIVITY

The specification given in Section III-C states that the probability of generating a warning message in response to the [illegible] of M = 5 or fewer vehicles within an interval of 5 minutes shall be less than Q = 0.1. An important question in this regard is the extent to which the event descriptions extracted from the [illegible] signatures provide means for distinguishing between threatening and threats. We will not address ourselves to this question here, [illegible] to note that even the most rudimentary event-classification [illegible] should be of considerable value in suppressing unwanted warning messages in response to nonbelligerent activity.

A second question is whether the system could ever generate a warning message in response to the launching of 5 vehicles when the warning threshold is set above 5, e.g., $N_W = 11$. The answer is clearly no, if a single launching at most can give rise to one event detection by the system. As has been noted in the previous section, however, each event may give rise to more than one detected signature, because the event may generate signatures in more than one sensor in the system array. Thus, unless multiple signatures from a single event can be perfectly associated, the possibility exists that the launching of fewer than N_W vehicles can lead to an event count that exceeds N_W and causes the generation of a warning message.

The analysis supporting this discussion is presented in Appendix B. The assumption is made that each event leads to a response (in the form of either no detection or one or more event detections) that is statistically independent of the responses made to the other concurrent events. This response is characterized by a probability vector

$$\vec{P}_E = \left(P_E^{(0)}, P_E^{(1)}, P_E^{(2)}, \ldots\right) \tag{13}$$

where $P_E^{(k)}$ is the probability that the event leads to k event detections. The probability that the event is detected at least once is now

$$P_E = 1 - P_E^{(0)} = \sum_{k=1}^{\infty} P_E^{(k)} \tag{14}$$

It will be recognized that Eq. 8 (for the probability of warning) is conservative, in the sense that it reflects at most a single event detection in response to an event.

To illustrate the sensitivity of Q to the probability of double and higher order detections, a few numerical examples will be presented. In all cases, it is assumed that $N_W = 11$ and that 5 vehicles were launched.

Case I

k	$P_E^{(k)}$	
0	0.25	$P_E = 0.75$
1	0.40	
2	0.20	$Q = 0.06$
3	0.10	
4	0.05	

Here the association of multiple signatures is fairly good; for 53+ percent of the events that are detected, the aggregation process yields a single event detection for each event. The requirement on Q is satisfied.

Case II

k	$P_E^{(k)}$	
0	0.2	$P_E = 0.80$
1	0.2	
2	0.2	$Q = 0.44$
3	0.2	
4	0.2	

Here the association of multiple signatures is somewhat poorer than in Case 1; for only 25 percent of the events that are detected, the aggregation process yields a single event detection in response to a single event. The requirement on Q is not met.

Case III

k	$P_E^{(k)}$	
0	0.25	$P_E = 0.75$
1	0.05	
2	0.10	$Q = 0.69$
3	0.20	
4	0.40	

In this case, the association is rather poor; in 53+ percent of the events that are detected, the aggregation process yields 4 event detections in response to a single event. The odds are better than two to one that the system will generate a warning message.

The implication of these results is that the specification constraint on the warning system response to nonbelligerent activity imposes a restriction on the design of the system. It is more or less obvious that the probability of event detection can be increased by increasing the number of sensors in the warning system sensor array, because the number of signatures delivered by the array in response to an event can be increased. It is equally more or less obvious, however, that the difficulty of achieving perfect aggregation of multiple signatures from a single event (so as to yield a single event detection) increases with the number of signatures that are made available. If (as a result of an attempt to increase the probability of event detection) the aggregation performance is moderately degraded, the effect can be to cause the system to generate warning messages in response to nonbelligerent activity with an undesirably high probability. Thus, concomitant with a determination of the number and coverage of sensors in the sensor array, there must be a determination of the capabilities of the system for aggregating multiple signatures provided by the array.

V. IMPLICATIONS FOR THE WARNING COMPLEX

A. THE INTEGRATION CONCEPT

The discussion in Section III-B indicated that there are a large variety of ways in which the outputs of several warning systems in a complex can be combined. In this section, we shall consider a most rudimentary scheme, that of basing decisions by the complex solely on the occurrence of warning messages from the individual systems. The descriptive content of the messages is ignored, insofar as the decision by the complex to transmit a warning message to the user is concerned. It will be noted, however, that matching the descriptive content of messages from different warning systems is a very powerful technique for suppressing false-warning messages from the complex and is practically essential if the complex is to provide a meaningful description of an impending attack.

For concreteness, it will be assumed that three statistically independent warning systems provide the inputs used by the complex in deciding whether to transmit a warning message to the user. The sequence of false-warning messages generated by each system is assumed to be governed by the Poisson distribution. Three possible decision rules for use by the complex will be considered.

- 1/3: The complex transmits a warning message in response to a warning message from one or more of the three systems.
- 2/3: The complex transmits a warning message when a message from one system is corroborated by a message from one or both of the other systems.
- 3/3: The complex transmits a warning message when a message from one system is corroborated by messages from both of the other systems.

For the second and third rules to be meaningful, it is necessary to specify what is meant by corroboration. Specifically, a warning message from one of the three systems will be said to be corroborated if a warning message arrives from another of the systems within a time interval (specified relative to the time of reception of the first warning message) of specified duration, T_C minutes. The specification of T_C depends on *a priori* knowledge of the uncertainty in the time of reception of the corroborating message relative to the time of reception of the original message. The corroboration interval itself may begin as soon as the first warning message is received, or it may commence after a predetermined delay.

The criteria used by the individual systems for generating initial warning messages need not be the same as the criteria for generating corroborating warning messages. Practically speaking, the user of warning messages will undoubtedly take certain actions in response to a message from a single warning system, and more significant actions in response to one or more corroborations of the initial message.*

The possibility also exists for a check-back corroboration rule. Under this rule, reception of a warning message from one system elicits a search for prior activity observed by another system that would normally have delivered its warning message earlier than the system that actually sent the warning. This permits corroboration from event detections that were insufficient to cause generation of a warning message by the system that is normally earlier.

B. DETECTION AND FALSE-WARNING RATE PERFORMANCE

Under the assumptions that have been stated, the probability that the complex will generate a warning message is given by

$$P_W^{(c)}(1/3) = 1 - \prod_{i=1}^{3} \left[1 - P_W^{(i)}\right] \qquad (15)$$

*In this regard, all three of the decision rules being considered are therefore relevant, different weights being attached to messages generated by the complex under different rules.

under the 1/3 rule, where $P_W^{(i)}$ is the probability that the i^{th} system will generate a warning message. If P_W is the same for all threee systems, then

$$P_W^{(c)} = 1 - \left(1 - P_W\right)^3 \tag{16}$$

For the 2/3 rule,

$$\begin{aligned} P_W^{(c)}(2/3) = {} & P_W^{(1)}\left[1 - \left(1 - P_C^{(2)}\right)\left(1 - P_C^{(3)}\right)\right] \\ & + P_W^{(2)}\left[1 - \left(1 - P_C^{(1)}\right)\left(1 - P_C^{(3)}\right)\right] \\ & + P_W^{(3)}\left[1 - \left(1 - P_C^{(1)}\right)\left(1 - P_C^{(2)}\right)\right] \\ & + P_W^{(1)}\,P_W^{(2)}\,P_W^{(3)} - P_W^{(1)}\,P_W^{(2)} - P_W^{(1)}\,P_W^{(3)} - P_W^{(2)}\,P_W^{(3)} \end{aligned} \tag{17}$$

where $P_C^{(i)}$ is the probability that the i^{th} system will generate a corroborating warning message. If tne three systems yield equal values for P_W and P_C, then Eq. 17 simplifies to

$$P_W^{(c)}(2/3) = P_W^3 - 3P_W^2 + 3P_W\left[1 - \left(1 - P_C\right)^2\right] \tag{18}$$

For the 3/3 rule,

$$\begin{aligned} P_W^{(c)}(3/3) = {} & P_W^{(1)}\,P_W^{(2)}\,P_W^{(3)} + P_W^{(1)}\,P_C^{(2)}\,P_C^{(3)} + P_C^{(1)}\,P_W^{(2)}\,P_C^{(3)} \\ & + P_C^{(1)}\,P_C^{(2)}\,P_W^{(3)} - P_W^{(1)}\,P_W^{(2)}\,P_C^{(3)} - P_W^{(1)}\,P_C^{(2)}\,P_W^{(3)} - P_C^{(1)}\,P_W^{(2)}\,P_W^{(3)} \end{aligned} \tag{19}$$

Again, if the three systems yield identical values for P_W and P_C,

$$P_W^{(c)}(3/3) = P_W^3 + 3P_W\,P_C^2 - 3P_W^2\,P_C \tag{20}$$

Table 3 gives values for $P_W^{(c)}$ (1/3), $P_W^{(c)}$ (2/3), and $P_W^{(c)}$ (3/3), as given by Eqs. 16, 18, and 20, as functions of P_W, under the assumption that $P_C = P_W$, which is equivalent to saying that identical criteria are employed for generating initial and corroborating warning messages. It can be seen that the triple redundancy provided by the 1/3 rule yields values for the probability of warning by the complex which, for $P_W = 0.5$, exceed those obtained using the 3/3 rule with $P_W = 0.95$. If the single-system specification of $P_W = 0.95$ is met, however, then the value of $P_W^{(c)}$ that is obtained using the 3/3 rule may be satisfactorily high.

TABLE 3. PROBABILITY OF WARNING FROM THE COMPLEX AS A FUNCTION OF PROBABILITY OF WARNING FROM A SINGLE SYSTEM (THREE-SYSTEM COMPLEX)

P_W	$P_W^{(c)}$ (1/3)	$P_W^{(c)}$ (2/3)	$P_W^{(c)}$ (3/3)
0.95	0.9999	0.9927	0.8574
0.90	0.9990	0.9720	0.7290
0.85	0.9966	0.9392	0.6141
0.80	0.9920	0.8960	0.5120
0.75	0.9844	0.8437	0.4219
0.70	0.9730	0.7840	0.3430
0.60	0.9360	0.6480	0.2160
0.50	0.8750	0.5000	0.1250

Although the 2/3 rule does not yield values of $P_W^{(c)}$ as high as those obtained with the 1/3 rule, it does give values that are quite respectable over a range of attainable values for P_W. Corroboration in effect degrades the probability-of-warning performance of the complex, relative to that achievable using the 1/3 rule, but the false-warning rate associated with the 2/3 rule is vastly superior to the false-warning rate associated with a 1/3 complex.

The false-warning rate associated with a complex using the 1/3 rule is given by

$$R_W^{(c)}(1/3) = \sum_{i=1}^{3} R_W^{(i)} \tag{21}$$

For the 2/3 and 3/3 rules, it is convenient to define a false-corroboration probability, $Q_C^{(i)}$, which is the probability that the i^{th} system will generate a false corroborative warning message during the corroboration time interval (of duration T_C).

For the complex using the 2/3 rule, the false-warning rate is given by

$$\begin{aligned} R_W^{(c)}(2/3) = {} & R_W^{(1)} \left[Q_C^{(2)} + Q_C^{(3)} - Q_C^{(2)} Q_C^{(3)}\right] \\ & + R_W^{(2)} \left[Q_C^{(1)} + Q_C^{(3)} - Q_C^{(1)} Q_C^{(3)}\right] \\ & + R_W^{(3)} \left[Q_C^{(1)} + Q_C^{(2)} - Q_C^{(1)} Q_C^{(2)}\right] \end{aligned} \tag{22}$$

which reduces to

$$R_W^{(c)}(2/3) = 3R_W \left[2 Q_C - Q_C^2\right] \tag{23}$$

when the systems are statistically identical.

The false-warning rate of a complex using the 3/3 rule is given by

$$R_W^{(c)}(3/3) = R_W^{(1)} Q_C^{(2)} Q_C^{(3)} + R_W^{(2)} Q_C^{(1)} Q_C^{(3)} + R_W^{(3)} Q_C^{(1)} Q_C^{(2)} \tag{24}$$

which simplifies, when the three systems are statistically identical, to

$$R_W^{(c)}(3/3) = 3R_W Q_C^2 \tag{25}$$

The problem of calculating Q_C is nontrivial, involving what are known as clustering statistics. For example, suppose that for corroboration purposes the configuration of Figure 3 is used, with a threshold N_C. Then Q_C is the probability that N_C false-event detections occur during an interval of duration T_A, at some time during the corroboration interval T_C. First, let $C(k; m, T_A/T_C)$ be the probability that if m false-event detections occur within the interval T_C, k or more of them fall within an interval of duration T_A. The computation of this probability is discussed in Appendix C. Under the assumption that false-event detections are governed by the Poisson distribution, the probability of a false corroborative warning message is given by

$$Q_C = \sum_{m=N_C}^{\infty} \frac{(R_E T_C)^m}{m!} \exp(-R_E T_C) C(N_C, m; T_A/T_C) \quad (26)$$

where R_E, as before, is the false-event detection rate. It will be recalled that for T_A = 5 minutes, R_E = 20.6 per hour, and N_W = 11, a system false-warning rate, R_W, of 2 per year is obtained. Assuming T_C = 10 minutes,* the following values are obtained in Appendix C for Q_C:

$$Q_C = 0.00265 \qquad \text{for } N_C = 8$$

and

$$Q_C = 0.000019 \qquad \text{for } N_C = N_W = 11$$

Using these results, a comparison between the average false-warning rates for complexes using the 1/3, 2/3, and 3/3 rules can be made by means of Eqs. 21, 23, and 25:

$$R_W^{(c)}(1/3) = 6 \text{ per year} = 1/2 \text{ months}$$

*This value is probably fairly conservative.

$N_C = 8$

$R_W^{(c)}(2/3) = 0.032$ per year $=$ 1/31 years

$R_W^{(c)}(3/3) = 4.2 \times 10^{-5}$ per year $=$ 1/24,000 years

$N_C = 11$

$R_W^{(c)}(2/3) = 2.4 \times 10^{-4}$ per year $=$ 1/4200 years

$R_W^{(c)}(3/3) = 1.2 \times 10^{-9}$ per year $=$ 1/800 million years

This comparison shows the rather profound effect that the corroboration requirements of the 2/3 and 3/3 rules have on the false-warning rate from the complex.

In addition, it seems apparent that the criteria used by the individual warning systems for corroboration need not be as stringent as the criteria for initial warnings. It should be noted in this regard that the use of a corroboration threshold (N_C) that is lower than the warning threshold was not reflected in the warning probability calculations presented in Table 3, and the results presented there for complexes using the 2/3 and 3/3 decision rules must therefore be regarded as conservative. For example, the first entry in Table 3 ($P_W = 0.95$) is obtained for $N_A = 20$, $N_W = 11$, and $P_E = 0.7$. For these parameter values, the corresponding corroboration probability, P_C, is 0.9987 (instead of 0.95) if $N_C = 8$, and the probability of warning for a complex using the 3/3 rule is 0.9960 instead of 0.8574, as given in Table 3.

It should be noted that the effect of multiple false-event detections of the variety discussed in Section IV-B will be to increase the values obtained for Q_C, perhaps dramatically.

C. WARNING TIME CONSIDERATIONS

The discussions presented here are quite primitive and are based on the elementary remark that if the warning complex employs a decision rule requiring corroboration, such as the 2/3 and 3/3 rules that have been considered above, then the complex must wait until the corroboration requirements have been met before it can transmit a warning message to the user. Conversely, a complex that does not require corroboration can deliver a warning message as soon as it receives a warning message from the first system to detect an attack. This fairly obvious point may have some bearing on the selection of warning systems in general; it certainly influences the design of individual warning systems in the complex. For example, if, in a two-system complex, one of the systems can be expected to deliver a warning message with an average warning time of 20 minutes (for a particular threat class), and if such a message must be corroborated by a message from a second warning system that can be expected to respond with an average warning time of 10 minutes, then there is little justification for improving the warning time of the first system to (say) 25 minutes, because of the obvious point that the average warning time provided by the complex would still be 10 minutes.

To provide an analytical assessment of the effects of corroboration requirements, a comparison will be made between a three-system warning complex using the 1/3 decision rule and a similar complex using the 3/3 rule. It will be assumed that the warning times provided by the individual warning systems are statistically independent and identically distributed. Specifically, it will be assumed that each warning system delivers a warning message in response to the onset of an attack with probability P_w (as before), and that if such a message is delivered, it will arrive at the complex at a random time t_w, which is uniformly distributed between T_{min} and T_{max}, the former time being the minimum warning time provided by the system (given that it generates a warning message at all), and the latter being the maximum warning time that the system is capable of providing.

The complex using the 3/3 rule cannot generate a warning un l it has received a warning message and two corroborations. Denot the warning time provided by the i^{th} system by $t_w^{(i)}$, and let

$$u_i = \frac{t_w^{(i)} - T_{min}}{T_{max} - T_{min}} \qquad i = 1,2,3 \qquad (\ )$$

Under the foregoing assumptions, u_1, u_2, and u_3 are statistically n-dependent and uniformly distributed over the interval (0, 1). Th probability that the smallest of u_i is less than u ($0 \le u \le 1$) is hen given by

$$\text{Prob}\left\{u_{min} < u\right\} = 1 - (1-u)^3 \qquad (\ )$$

This result corresponds to the probability distribution of the wa ing time afforded by the last of the three messages to arrive at the m-plex, which in turn is essentially the warning time provided by t complex (with appropriate adjustments in T_{min} and T_{max}). That is the probability distribution of the warning time provided by a co - plex using the 3/3 rule is given by

$$\text{Pro}\left\{t_w^{(c)}(3/3) < t\right\} =$$

$$3\left(\frac{t - T_{min}}{T_{max} - T_{min}}\right) - 3\left(\frac{t - T_{min}}{T_{max} - T_{min}}\right)^2 + \left(\frac{t - T_{min}}{T_{max} - T_{min}}\right)^3 (\ )$$

for $T_{min} \le t \le T_{max}$.

It is then a fairly straightforward matter to calculate the mean ing time provided by the 3/3 complex:

$$\bar{t}_w^{(c)}(3/3) = E\left\{t_w^{(c)}(3/3)\right\} = T_{min} + (1/4)\left(T_{max} - T_{min}\right) (\ )$$

A similar computation can be made for the warning complex using the 1/3 rule, except that the procedure is somewhat more complex. Let

$$P_1 = \frac{3P_W\left(1 - P_W\right)^2}{1 - \left(1 - P_W\right)^3} \tag{31}$$

$$P_2 = \frac{3P_W^2\left(1 - P_W\right)}{1 - \left(1 - P_W\right)^3} \tag{32}$$

$$P_3 = \frac{P_W^3}{1 - \left(1 - P_W\right)^3} \tag{33}$$

It will be recognized that P_k is the conditional probability that, given that at least one system transmitted a warning message to the complex, exactly k did so. It is next possible to compute the conditional probability distribution of the warning time provided by the complex, given that exactly k systems generated warning messages. The result obtained is that

$$\text{Prob}\left\{t_w^{(c)}(1/3) < t \,|k\right\} = \left(\frac{t - T_{min}}{T_{max} - T_{min}}\right)^k \tag{34}$$

and the overall probability distribution for the warning time provided by the 1/3 complex is:

$$\text{Prob}\left\{t_w^{(c)}(1/3) < t\right\} = \sum_{k=1}^{3} P_k \left(\frac{t - T_{min}}{T_{max} - T_{min}}\right)^k \tag{35}$$

The average warning time provided by the 1/3 complex is readily calculated to be

$$T_W^{(c)}(1/3) = E\left\{t_w^{(c)}(1/3)\right\} \tag{36}$$

$$= T_{min} + \left(T_{max} - T_{min}\right)\left[(P_1/2) + (2P_2/3) + (3P_3/4)\right]$$

For $P_W = 0.95$, Eqs. 31-33 and 37 yield

$$T_W^{(c)}(1/3) = T_{min} + 0.737\ (T_{max} - T_{min}) \qquad (37)$$

Thus, if $T_{max} - T_{min}$ is 5 minutes, the 1/3 complex will provide an average warning time that is about 2.4 minutes greater than the average warning time provided by the 3/3 complex, under the many assumptions that have been made.

Computations for a warning complex using the 2/3 decision rule will not be presented, but it can be stated that the warning time advantage enjoyed by the 1/3 complex over the 2/3 complex is considerably less than the advantage over the 3/3 complex. Moreover, it should be noted that the time required to obtain corroborative warning messages may be less than that required to generate initial warning messages, if the criteria for corroboration are less stringent than those for initial warnings.

The probability distribution assumed for the warning time provided by an individual system is, of course, purely hypothetical. The actual distribution will depend not only on the sensor characteristics and communication and processing times, but also on the attack scenario.

VI. SUMMARY AND CONCLUSIONS

We have reviewed several conceptual questions dealing with warning system structures and organization, and means for combining the outputs of multiple warning systems into an aggregated response for a warning complex. The following points should be stressed:

1. The necessity for multiple and diverse warning systems to provide adequate spatial and temporal coverage of a variety of present and future threats
2. The necessity for correlating the outputs from such systems in order to provide effective and accurate descriptions of impending attacks and to minimize false warnings and their consequences
3. The inherent (but not necessarily insurmountable) difficulty of correlating outputs from diverse types of warning systems

In these discussions, several important questions have not been discussed. First, the role of human participants in the warning decision process has been completely ignored. The position taken has been one of examining the mechanistic possibilities for warning systems and the warning complex, recognizing that human participants can always be added to provide procedural modifications, criteria adaptation, and other supervisory functions, and to accommodate situations that were not foreseen when the equipment was designed. This viewpoint may be somewhat contrary to the generally accepted philosophy.

Second, the need for adaptive implementation and operational concepts has not received proper stress. Changes in the environment of the sensing elements, changes in the structure of the warning complex, and new threat classes (which may not be immediately identifiable as threats) impose a requirement for flexibility and selectivity in system capabilities and procedures.

Third, the interactions between tactical warning, strategic offensive forces, and strategic defenses have not been properly assessed. There is an important trade-off between warning capabilities, offensive force requirements, and defensive requirements that should be examined in depth, with the aim of specifying a realistic balance of resources allocated to these components of deterrence.

In an analytical vein, we have attempted to illustrate some of the problems of determining the performance of a warning system and a warning complex. The consequences of imperfect association of multiple sensor responses to false and real events have been examined, the conclusion being that while provision of redundant sensing capabilities can provide better coverage and improved attack detection probabilities, the advantages may be offset to a significant degree by false-warning rate increases and undesirable warnings generated in response to nonbelligerent activity. On the positive side, it has been found that even the correlation of warning system outputs at the grossest possible level is capable of yielding acceptable probabilities of warning from the complex, while, at the same time, the false-warning rates obtained can be easily reduced to an acceptably low rate, e.g., one false warning per several decades or even many millenia. Of the configurations examined, the complex using the 2/3 decision rule appears to offer the best combination of detection, false-warning rate, and reaction time performance, although such an inference should be clearly recognized as being sensitive to the many assumptions made during the analyses.

The fact that such performance can be achieved with such apparent ease is not to be construed as obviating the need for higher order techniques for integrating warning system outputs. In the final analysis, the output of the warning complex must be used by human decision makers. The bald statement that an attack has been detected and corroborated by two out of three warning systems is not likely to be as effective as such a statement accompanied by an integrated description of the attack, in which the description provided by one of the systems is demonstrated to be consistent with and reinforced by the description

provided by the other. Moreover, in addition to corroboration in detail, the aggregate description of the attack is likely to be significantly more comprehensive and accurate than any of the individual descriptions and is therefore likely to be far more useful to the final decision maker.

APPENDIX A

COUNTING STATISTICS FOR COMPOUND POISSON PROCESSES

It is supposed that false-event detections can occur singly, in pairs, or in groups of three, as described in Section IV-B. The sequence of singlets, the sequence of doublets, and the sequence of triplets are assumed to be mutually statistically independent, and each sequence is assumed to be governed by the Poisson distribution. The determination of the false-warning rate associated with this combination of false-event-detection sequences will be broken down into three mutually exclusive cases.

1. Suppose that a <u>single</u> false-event detection has just been delivered to the counter; then a warning message will be generated in response to that detection if and only if exactly $N_W - 1$ false-event detections occurred in the previous T_A minutes. The average rate of occurrence of singlets is denoted by $R_E^{(1)}$; thus, the rate of false-warning messages that are triggered by singlets is

$$R_W^{(1)} = R_E^{(1)} \times P\,(N_W - 1,\ T_A) \qquad \text{(A-1)}$$

where P (n, T) is the probability of exactly n false-event detections in T minutes.

2. Suppose that a <u>double</u> false-event detection has just been delivered to the counter; then a warning message will be generated if either $N_W - 2$ or $N_W - 1$ false event detections occurred in the previous T_A minutes. In the first instance, the warning message will be triggered by the first detection of the pair. The average rate of occurrence of doublets is denoted by $R_E^{(2)}$; thus, the rate of false-warning messages that are triggered by doublets is

$$R_W^{(2)} = R_E^{(2)}\left[P\left(N_W - 1, T_A\right) + P\left(N_W - 2, T_A\right)\right] \qquad \text{(A-2)}$$

3. Finally, suppose that a <u>triple</u> false-event detection has been delivered to the counter. A warning message will be generated if $N_W - 3$, $N_W - 2$, or $N_W - 1$ false-event detections occurred in the previous T_A minutes. In the first instance, the warning message will be triggered by the first of the three false-event detections; and in the second, the warning message will be triggered by the second of the three false-event detections in the triplet. The average rate of occurrence of triplets is denoted by $R_E^{(3)}$; thus, the rate of false-warning messages that are triggered by triplets is

$$R_W^{(3)} = R_E^{(3)}\left[P\left(N_W - 1, T_A\right) + P\left(N_W - 2, T_A\right) + P\left(N_W - 3, T_A\right)\right] \qquad \text{(A-3)}$$

Because the three cases just considered are mutually exclusive, the overall false-warning rate can be calculated by addition.

$$R_W = R_W^{(1)} + R_W^{(2)} + R_W^{(3)} \qquad \text{(A-4)}$$

Combining these results yields

$$\begin{aligned} R_W = {} & P\left(N_W - 1, T_A\right)\left[R_E^{(1)} + R_E^{(2)} + R_E^{(3)}\right] \\ & + P\left(N_W - 2, T_A\right)\left[R_E^{(2)} + R_E^{(3)}\right] \\ & + P\left(N_W - 3, T_A\right)\left[R_E^{(3)}\right] \end{aligned} \qquad \text{(A-5)}$$

The problem is therefore reduced to one of calculating P (n, T), given the rates $R_E^{(1)}$, $R_E^{(2)}$, and $R_E^{(3)}$. Analytically, this expression is given by

$$P(n, T) = \exp\left[- R_E^{(1)} T - R_E^{(2)} T - R_E^{(3)} T\right] \times \sum\sum\sum \frac{\left[R_E^{(1)} T\right]^{k_1}\left[R_E^{(2)} T\right]^{k_2}\left[R_E^{(3)} T\right]^{k_3}}{k_1!\ k_2!\ k_3!} \tag{A-6}$$

The sum in Eq. A-6 is taken over the non-negative values of k_1, k_2, and k_3, for which

$$k_1 + 2k_2 + 3k_3 = n \tag{A-7}$$

An easy way to compute P (n, T) is to allow k_1, $2k_2$, and $3k_3$ to run from zero to the maximum value of n (i.e., $N_W - 1$); the summand of Eq. A-6 is computed only when the value of n obtained from Eq. A-7 is of interest. The computed summands are then sorted according to Eq. A-7, and added to sums corresponding to the values that are of interest. After completion of this process, the sums are multiplied by the exponential indicated in Eq. A-6 to obtain a table of P (n, T).

APPENDIX B

MULTINOMIAL DISTRIBUTIONS

A single event can lead to 0, 1, 2, ... event detection . It is assumed that the number of event detections obtained by the system from an event is statistically independent of the number obtained from every other event and that the statistical characterizations of the events are identical. Let $P_E^{(k)}$ denote the probability that an event causes the system to obtain k event detections following aggregation. Then if M events occur, the probability that the system will obtain k_1 unaggregated detections of the first event, k_2 unaggregated detections from the second event, ..., and k_M unaggregated detections from the M^{th} event is simply

$$\text{Prob}\left\{k_1, k_2, \ldots, k_M\right\} = \prod_{i=1}^{M} P_E^{(k_i)} \qquad \text{(B-1)}$$

The probability Q of generation of a warning message is obtained by summing this expression over all values of $k_1, k_2, \ldots, k_M$ for which

$$k_1 + k_2 + \ldots + k_M \geq N_W \qquad \text{(B-2)}$$

That is,

$$Q = \sum_{k_1}\sum_{k_2}\cdots\sum_{k_M} \prod_{i=1}^{M} P_E^{(k_i)} \qquad \text{(B-3)}$$

where each index ranges from 0 to K, which is to be chosen so that

$$P_E^{(k)} = 0 \qquad \text{for all } k > K$$

and the summation is constrained by Eq. B-2.

Computation of Q is easily accomplished by a sorting technique akin to that mentioned in Appendix A. Alternatively, M numbers are allowed to run through all possible combinations (each number ranging from 0 to K); when Eq. B-2 is satisfied, the summand of Eq. B-3 is computed and added to a sum for Q.

APPENDIX C

CLUSTERING STATISTICS

Consider a random time sequence of events that is characterized by the Poisson distribution and by an average occurrence rate R_E. The probability that exactly m events occur in a time interval T_C is then given by

$$P(m) = \frac{(R_E T_C)^m}{m!} \exp(-R_E T_C) \qquad \text{(C-1)}$$

The problem to be discussed here is the determination of the probability $Q_C(N_C)$ that N_C or more such events will be clustered within some interval of duration T_A contained within T_C. This question is treated at length in a paper by Turner and Warren.* The approach is to determine in one way or another the probability $C(N_C, m; T_A/T_C)$ that, given exactly m events in the interval T_C, N_C or more will fall within some interval of duration T_A contained within T_C. Q_C is then obtained by multiplying $C(N_C, m; T_A/T_C)$ by the probability of exactly m events, P(m), and summing over those values of m for which $C(N_C, m; T_A/T_C)$ is nonzero:

$$Q_C = \sum_{m=N_C}^{\infty} P(m)\, C(N_C, m; T_A/T_C) \qquad \text{(C-2)}$$

which, using Eq. C-1, yields Eq. 26 of the text.

It is shown in the referenced paper that the stipulation that exactly m Poisson-distributed events have occurred in T_C is equivalent

*Robert P. Turner and Wayne P. Warren, Clustering Statistics of Uniformly Distributed Random Variables (U), Research Paper P-345 (Arlington, Va.: Institute for Defense Analyses, November, 1967) (SECRET). The results cited here are unclassified.

to a stipulation that the m events are uniformly distributed ove
Moreover, normalizing the time variable with respect to T_C leads
an equivalent statement of the problem of calculating C: given r andom
points uniformly distributed over the unit interval (0, 1), what
the probability that m or more of the events fall within an inte al
a ($= T_A/T_C$). Equivalently, one asks for the probability that the
smallest interval containing m points is less than a.

A general analytical expression for C(k, m; a) is not known, al-
though J.I. Naus has published a number of papers* containing spe al-
ized results. In general, C(k, m; a) is an m^{th} order polynomial
a; two known results are

$$C(m, m; a) = ma^{m-1} - (m-1)a^m \qquad (($$

which is the distribution of the range, and

$$C(m-1, m; a) = a^{m-2}\left[(1-2a)\, m\, (m-1) + a^2\, (m^2 - m + 2)\right]$$

$$+ \begin{cases} 0 & 0 \le a < 1/2 \\ (2a - 1)^m & 1/2 \le a \le 1 \end{cases} \qquad ((\ $$

In the referenced papers, a Monte Carlo technique was used t
estimating C(k, m; a) for $2 \le k < m$, $2 < m \le 15$. Comparison with
theoretical results indicated that the Monte Carlo estimates were
ali to two decimal places. The data obtained from that paper f
use in the example of the text are as follows:

---

*E.g., Joseph I. Naus, "The Distribution of the Size of the Maximum Cluster of Points on a Line," J. Amer. Stat. Assoc., pp. 532-38, 1965; "Some Probabilities, Expectations and Variances for the Size of Largest Clusters and Smallest Intervals," J. Amer. Stat. Assoc., pp. 1161-89, 1966.

m	C(8, m; 1/2)	C(11, m; 1/2)
8	0.0352*	0
9	0.1445*	0
10	0.35	0
11	0.55	0.00586*
12	0.74	0.0327*
13	0.88	0.11
14	0.98	0.22
15	1	0.36

These data were then combined using Eq. 26 of the text to yield the stated results for Q_C. It was assumed that C = 1 for m greater than 15.

*Computed using Eqs. C-3 and C-4.

UNCLASSIFIED
Security Classification

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1 ORIGINATING ACTIVITY (Corporate author) Institute for Defense Analyses	2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED 2b. GROUP --

3 REPORT TITLE

Remarks on Warning System Specifications and Structures

4 DESCRIPTIVE NOTES (Type of report and inclusive dates)

Research Paper P-534, August 1969

5 AUTHOR(S) (First name, middle initial, last name)

Robert D. Turner

6 REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
August 1969	53	--

8a. CONTRACT OR GRANT NO DAHC15 67 C 0011 b. PROJECT NO. A-5 c. d.	9a. ORIGINATOR'S REPORT NUMBER(S) P-534 9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report) None

10 DISTRIBUTION STATEMENT

This document has been approved for public release and sale; its distribution is unlimited.

11 SUPPLEMENTARY NOTES	12 SPONSORING MILITARY ACTIVITY
NA	Advanced Research Projects Agency Washington, D.C.

13 ABSTRACT

Several conceptual questions dealing with warning system structures are discussed, including techniques for correlating the outputs of multiple warning systems that compose a warning complex. The discussions are illustrated by a number of elementary analyses, dealing with warning probabilities, false-warning rates, and warning times. The mathematical calculations pertain to hypothetical automatic event detectors and automatic decision systems employing the outputs of such detectors. It is demonstrated that the use of multiple sensing elements in a warning system implies a need for proper association of redundantly detected signatures of real events and false signatures. The consequences of imperfect association are false-warning rates that can exceed by a considerable factor the rates that would be obtained assuming perfect association and when undesired warnings are generated in response to nonbelligerent activities.

DD FORM 1 NOV 68 1473

UNCLASSIFIED
Security Classification

Security Classification

14 KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT

Security Classification